https://doi.org/10.1038/s41567-022-01788-5

The flow from simulation to reality

Károly Zsolnai-Fehér

Fluid simulations today are remarkably realistic. In this Comment I discuss some of the most striking results from the past 20 years of computer graphics research that made this happen.

Growing up, I would often marvel at the smoke plumes ascending from a chimney and the water flows in the wake of a ship and assume that the underlying rules that describe them must be unfathomably complex. Later, as a student, I was struck by the deceptive simplicity of the Navier–Stokes equations, which — using only three terms codifying advection, pressure and diffusion — could describe waterfalls, waves around water droplets and turbulent smoke swirls.

Understandably, computer scientists have been eager to plug these equations into a computer and see the world come to life in their simulations. Unlike the computational fluid dynamics[1] literature, which aims for rigorous and accurate results, computer graphics research typically focuses on greater efficiency and artistic control, which are achievable with approximate solutions. These graphics solutions started appearing over 20 years ago[2,3], but hundreds of papers on this topic were still published in these 20 years — a testament to the complexity of the problem.

In the early days of computer graphics, simulations of liquids for animation and digital media involved evaluating the Navier–Stokes equations on a moving set of measurement points represented by particles, or on a stationary Cartesian grid. However, neither are suitable to the increasingly complex demands of the modern-day animation and media industries because of the huge computational cost. The particle technique requires too many particles to simulate detailed flows. The grid, on the other hand, suffers from the curse of dimensionality: for a detailed 3D scene, the grid points would have to be so finely laid out that even the most powerful supercomputer would be brought to its knees. Fortunately, computer graphics researchers have come far in the past 20 years and a supercomputer is no longer required. Here, I describe two of the community's ingenious ideas that can produce realistic animations of intricate fluid phenomena at a low computational cost.

One commonly used technique is spatial adaptivity[4]. Modern implementations[5] consider coarse grids for slow-moving waves and finer ones only in the regions where intricate details are likely to appear, for example around an area with turbulent motion (Fig. 1a, light and dark blue regions). This eases the computational burden a great deal with calmer regions requiring only a handful of grid points.

Another way to reduce computational cost is to simulate only (or mostly) the surface of liquids[6]. This idea is based on the key observation that although the majority of the simulation domain is often within the liquid, most of the fine details that are of interest for an animation appear on the surface. This approximation reduces the problem to two boundary integrals and eliminates the need to evaluate the full Navier–Stokes equations on a large volume. The extent of the simplification depends on the ratio of the volume to the surface of the fluid. Although it is helpful even in the case of a small droplet splashing into a thin sheet of water (Fig. 1b), the concept truly shines when simulating a large fluid body, such as an ocean.

However, simulating oceans poses challenges beyond computational cost. For example, diffuse effects such as bubbles and foam contribute a great deal to a realistic depiction. Unfortunately, adding these requires the introduction of surface tension calculations, which can get prohibitively expensive. However, leaning on the observation that bubbles and foam appear when air gets trapped within the fluid, one can limit the problem to wave crests, which can be easily identified by looking for regions where the curvature of the fluid geometry is high and locally convex[7]. This is, of course, an approximate solution, but its advantage is that foam and bubbles can be added to a finished, already existing simulation (Fig. 2a). And with this simple idea, the realism of a previously unconvincing simulation can be improved in a matter of minutes, even on a commodity computer.

Today, one need not despair even if surface-tension-based effects need to be included. A recent paper[8] offers three key realizations that show how much graphics research has progressed over the past 20 years. The first one is that, thanks to recent improvements in the efficiency of surface tension calculations, it is now possible to produce a realistic simulation of fluid phenomena as complex as cherries being dropped in a liquid. One can replicate the cherries being held up by buoyancy and capillary forces when dropped into water and getting submerged when dropped into milk (Fig. 2b). The second observation is that even these advanced surface-tension-based effects are no longer too costly to simulate. A modern solver can handle the interaction of liquids, membranes and solids at the same time, and thus obtain each

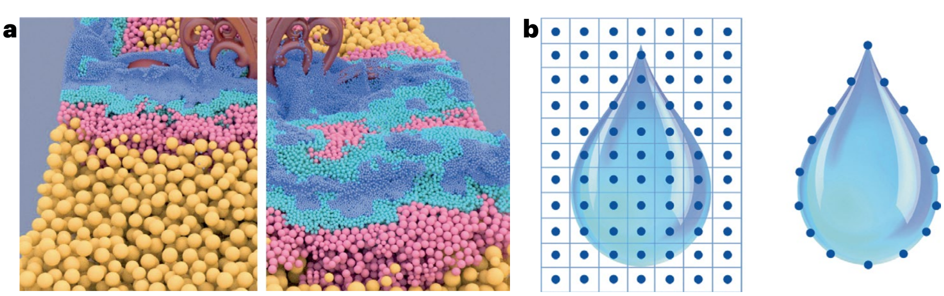

Fig. 1 | Two optimizations for fluid simulations. a, Spatial adaptivity is a simulation technique that allocates additional computational resources to areas with more complex flows, such as turbulent regions (light and dark blue). **b**, Simulation cost can be reduced by considering only the surface of liquids where most of the fine details appear.

CREDITS: A, REPRODUCED FROM REF. 5, ACM; B, COURTESY OF CHRIS WOJTAN AND EITAN GRINSPUN.

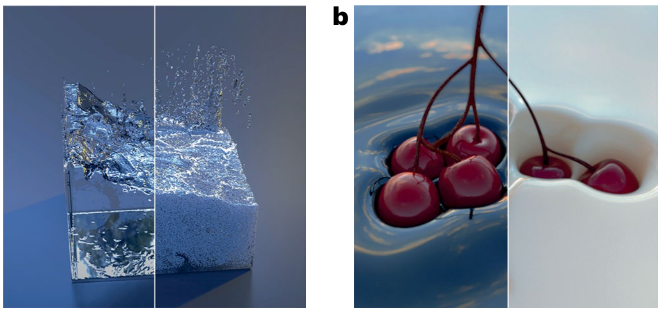

Fig. 2 | Advanced fluid simulation effects. a, The addition of diffuse effects, such as bubbles and foam, enhances the realism of a liquid simulation. **b**, Including advanced surface-tension-based effects can produce a realistic simulation of how buoyancy and capillary forces hold up or sink cherries dropped into different liquids, such as water and milk.

image in seconds. The third one is perhaps the most striking: when combined with a modern light transport simulator, the appearance of these simulations is now so convincing, they are almost indistinguishable from real-world photos.

Even with these improvements, it always seemed that in computer graphics, this realism is only there for looks – and I never expected these simulations to have any predictive power. But today, this prospect is becoming more and more likely. For decades, physics simulations for digital media were considered acceptable if they looked convincing to the human eye, and were nowhere near accurate enough for engineers to verify, for example, whether a new wind turbine design really does work correctly.

However, the computational cost of existing methods has decreased four-fold in just one year due to simpler, more efficient geometric approximation schemes that map more easily to existing graphics cards. With this, one can now simulate the airflow within a city block or create predictive wind tunnel tests for aircraft wing design with each second of animation taking only a few minutes to compute[9]. Simulations that are both real-time and predictive are within arm's reach – we might soon enter a world where an engineer is able to test new ideas in aircraft design every few minutes.

There are many more techniques that enable the simulation of intricate fluid phenomena such as the mesmerizing phenomenon of a ferrofluid climbing up a steel helix[10] or liquid–hair interactions[11]. It is also possible to extract the physical properties of a viscous material from a video recording of its dynamics[12].

These technical advances come at a price. More complex systems don't map well to existing hardware and their code base is more difficult to maintain and troubleshoot over time. Striking the right trade-off remains a key challenge when developing new simulation algorithms. However, this is also what makes this area a fertile ground for new ideas, where a small, but well-chosen compromise can introduce an order of magnitude increase in efficiency: those are the landmark papers in computer graphics.

On the other hand, neural-network-based learning approaches can generate increasingly convincing physics simulations more and more efficiently with each passing year[13]. Over time, they may even surpass conventional simulation methods. However, creating a unified system that can accommodate our appetite for realism and leave space for artistic directability, and do so efficiently enough to fulfil the requirements of modern artistic workflows, remains a challenge.

I would like to think that I have a vivid imagination, but after seeing all this progress I wonder what else we will be capable of 20 years and a few more papers down the line.

Károly Zsolnai-Fehér
Two Minute Papers, Pécs, Hungary.
e-mail: karoly@twominutepapers.com

Published online: 3 October 2022

References

1. Harlow, F. & Welch, J. E. *Phys. Fluids* **8**, 2182–2189 (1965).
2. Foster, N. & Metaxas, D. *Graph. Models Im. Proc.* **58**, 471–483 (1996).
3. Stam, J. Stable fluids. In *SIGGRAPH '99: Proc. 26th Conference on Computer Graphics and Interactive Techniques* 121–128 (ACM, 1999).
4. Berger, M. J. & Colella, P. *J. Comput. Phys.* **82**, 64–84 (1984).
5. Ando, R., Thürey, N. & Wojtan, C. *ACM Trans. Graph.* **32**, 103 (2013).
6. Da, F., Hahn, D., Batty, C., Wojtan, C. & Grinspun, E. *ACM Trans. Graph.* **35**, 78 (2016).
7. Ihmsen, M., Akinci, N., Akinci, G. & Teschner, M. *Visual Comput.* **28**, 669–677 (2012).
8. Ruan, L. et al. *ACM Trans. Graph.* **40**, 120 (2021).
9. Lyu, C., Li, W., Desbrun, M. & Liu, X. *ACM Trans. Graph.* **40**, 201 (2021).
10. Huang, L., Hädrich, T. & Michels, D. L. *ACM Trans. Graph.* **38**, 93 (2019).
11. Fei, Y., Maia, H. T., Batty, C., Zheng, C. & Grinspun, E. *ACM Trans. Graph.* **36**, 56 (2017).
12. Takahashi, T. & Lin, M. C. *ACM Trans. Graph.* **38**, 237 (2019).
13. Sanchez-Gonzalez, A. et al. Learning to simulate complex physics with graph networks. *Proc. 37th International Conference on Machine Learning* **119**, 8459–8468 (2020).

Acknowledgements
I thank C. Batty who graciously provided constructive feedback.

Competing interests
The author declares no competing interests.